

ORACLE

# The Future of Data and AI

---

**Tirthankar Lahiri**

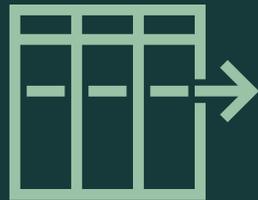
SVP Mission-Critical Data and AI Engines

Oracle



# The world of data was once simple

Traditional  
OLTP



On relational databases

Traditional  
Analytics



On relational databases

# Many new and specialized data technologies have been introduced

## New Types of Data



Relational



Documents



Spatial



Text

## New Types of Analytics



Data Warehouse



Data Lake



Graph



Streaming

## New Workload Types



AI



Geo-Distributed



IoT



Blockchain

# Modern apps are built by combining specialized databases

## New Types of Data



Relational



Documents



Spatial



Text

## New Types of Analytics



Data Warehouse



Data Lake



Graph



Streaming

## New Workload Types



AI



Geo-Distributed



IoT



Blockchain

# Complexity explodes further as we add AI to every component

## New Types of Data



Relational



Documents



Spatial



Text

## New Types of Analytics



Data Warehouse



Data Lake



Graph



Streaming

## New Workload Types



AI



Geo-Distributed



IoT



Blockchain

# Re-envisioning data management



Eliminate  
complexity



Improve  
synergy



Increase  
efficiency

# To achieve simplicity, engineer technology to work seamlessly together

## New Types of Data



Relational



Documents



Spatial



Text

## New Types of Analytics



Data Warehouse



Data Lake



Graph



Streaming

## New Workload Types



AI



Geo-Distributed



IoT



Blockchain

# Oracle engineers data management and app dev together to provide simplicity

## Engineered to Work Together



All data types and workloads



Hardware and software



OLTP, DW, and Data Lake



Full mission critical stack



Apps and data together

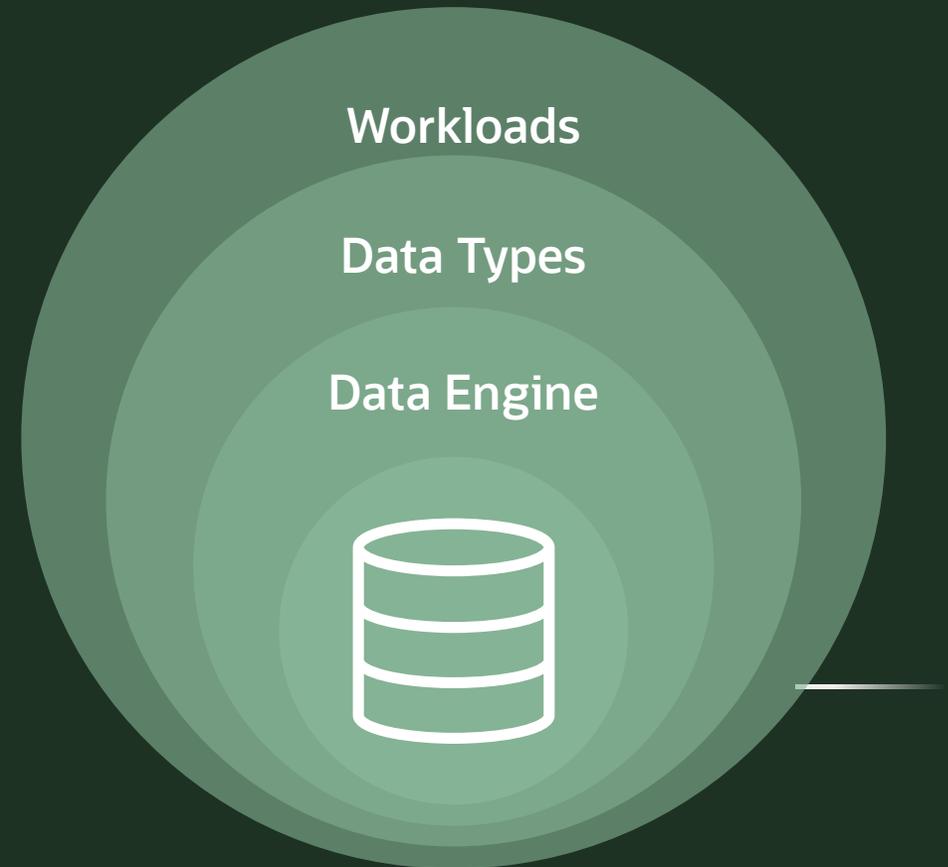


AI and app dev

# Oracle's Database Strategy

Engineer all data types and workloads to work together

Using a Converged Data Architecture



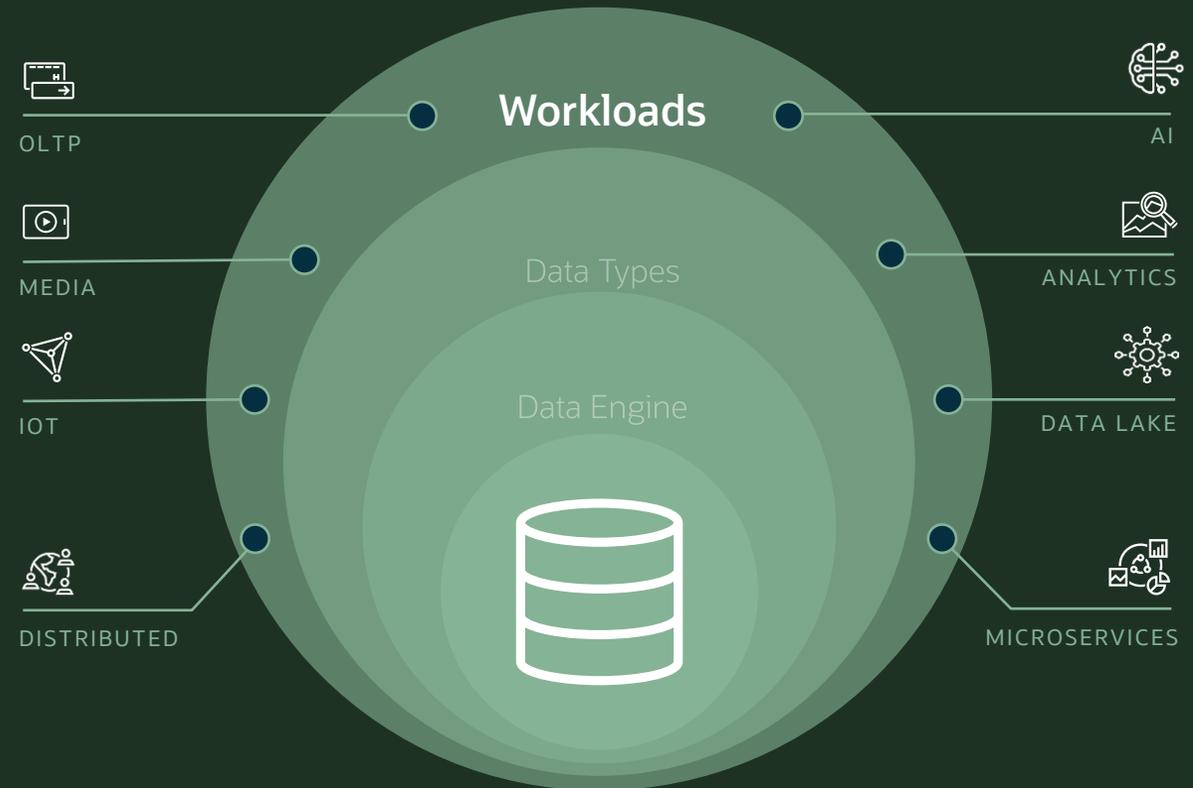
# Oracle's converged data architecture supports all key Workloads

Optimized performance for each workload

**Scales** to meet any need

Comprehensive support for apps using **mixed** workloads

**NEW:** Support for AI Vectors and Iceberg tables



All **Workloads** are Engineered to Work Together

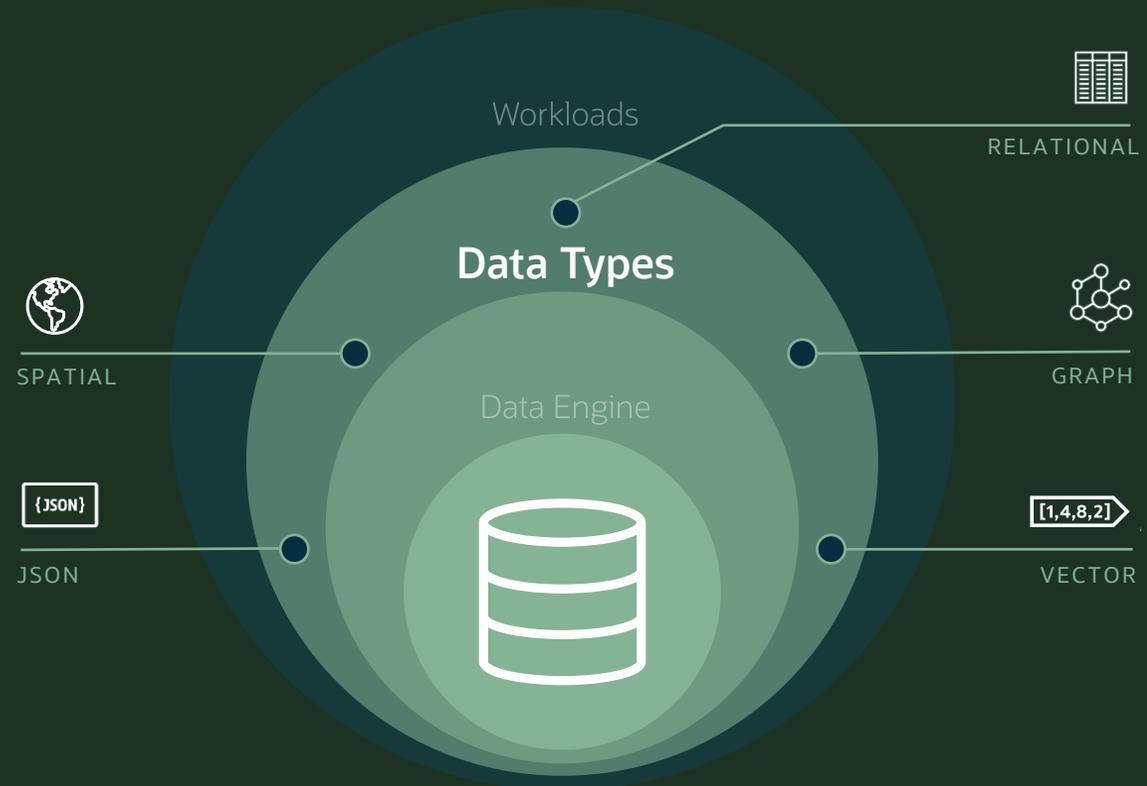
# Oracle's converged data architecture supports all key Data Types

Optimized for each type

And ideal for real-world apps that need a mix of multiple types

For a new data type or workload, add a SQL statement, not a DB

**NEW:** Ultra-fast JSON, MongoDB compatible API, and ISO SQL standard graph queries

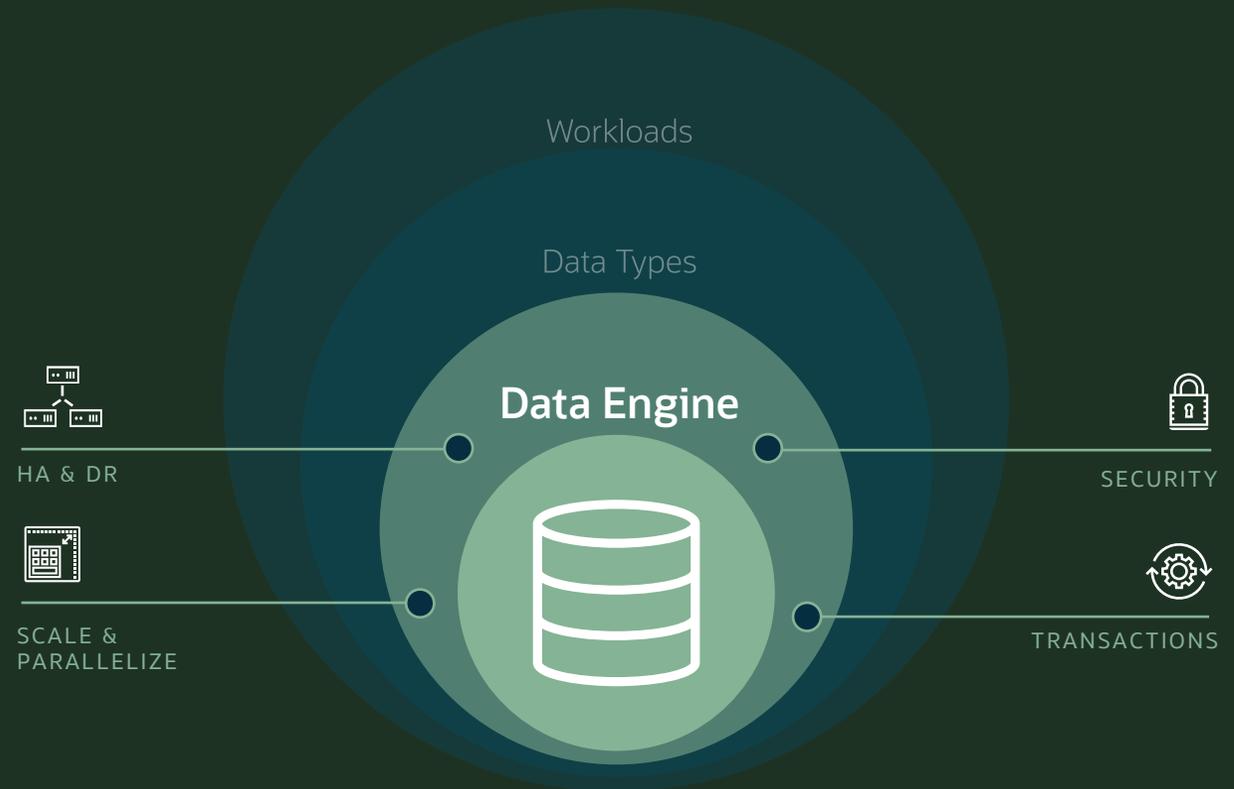


All **Data Types** are Engineered to Work Together

# Oracle's converged architecture runs on a common mission-critical Data Engine

Transparent atomicity, parallelism, high availability, disaster recovery, and security across all workloads and data types

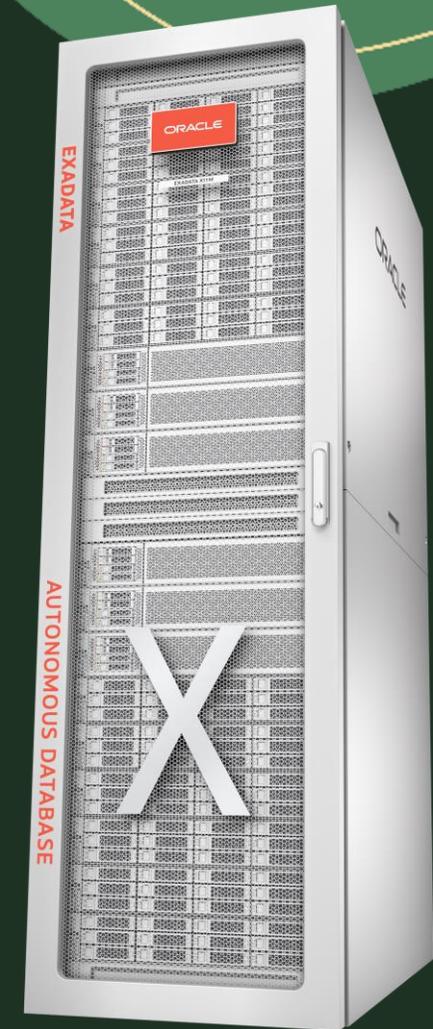
**NEW:** Lock-free updates and transactional microservices



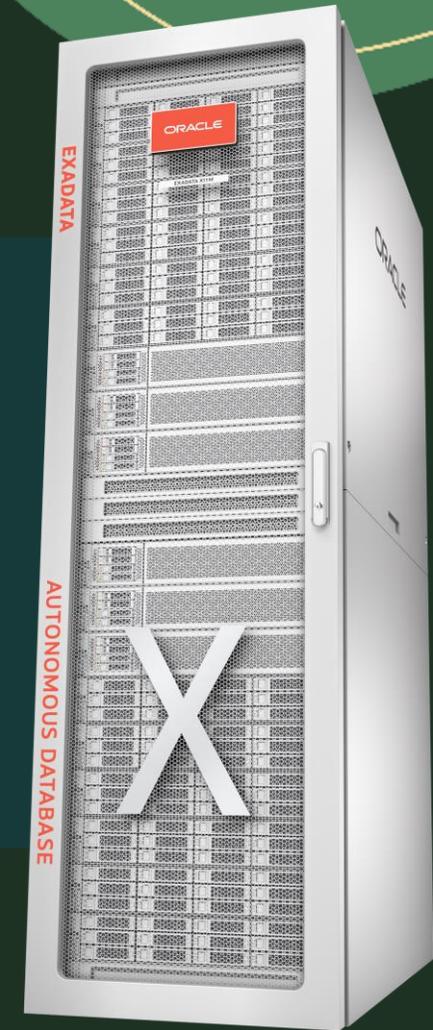
All **Data Operations** are Engineered to Work Together

# Oracle's Database Strategy

Engineer hardware and software together, everywhere



Exadata's unique smart storage, scale-out, and RDMA architecture delivers extreme performance and availability for all key workloads while reducing cost



**Software and Hardware** Engineered to Work Together

Thousands  
of global customers  
run on Exadata

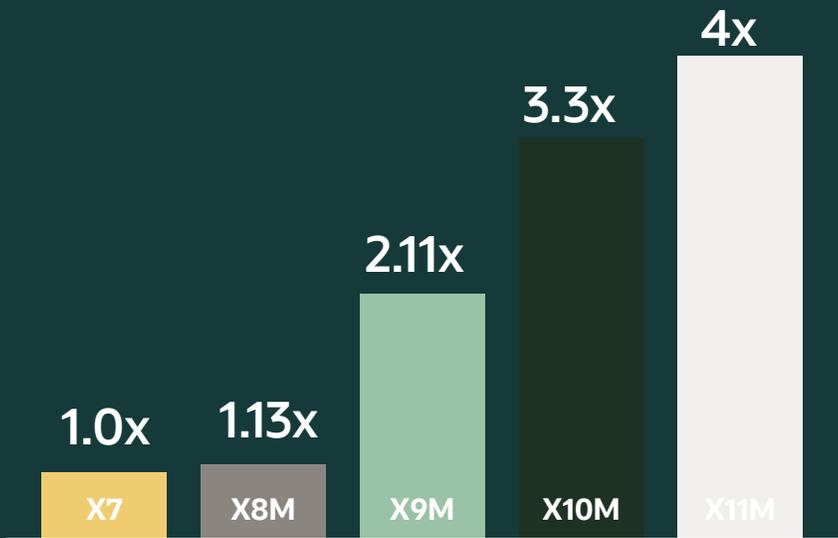
79% of Fortune  
Global 100 Run Exadata

58% Run Exadata Cloud



# Newest Exadata X11M – more performant and cost-effective than ever

up to **4x** Faster concurrent transaction throughput



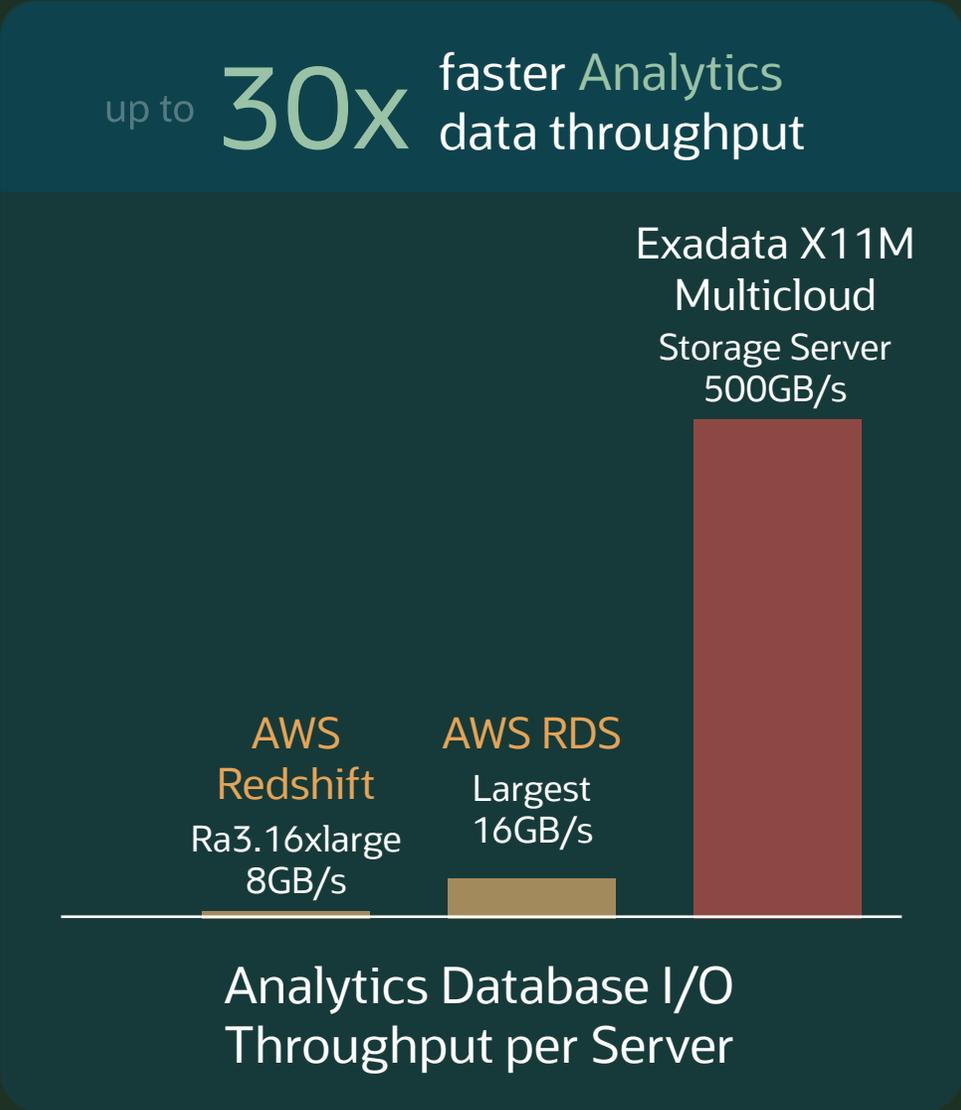
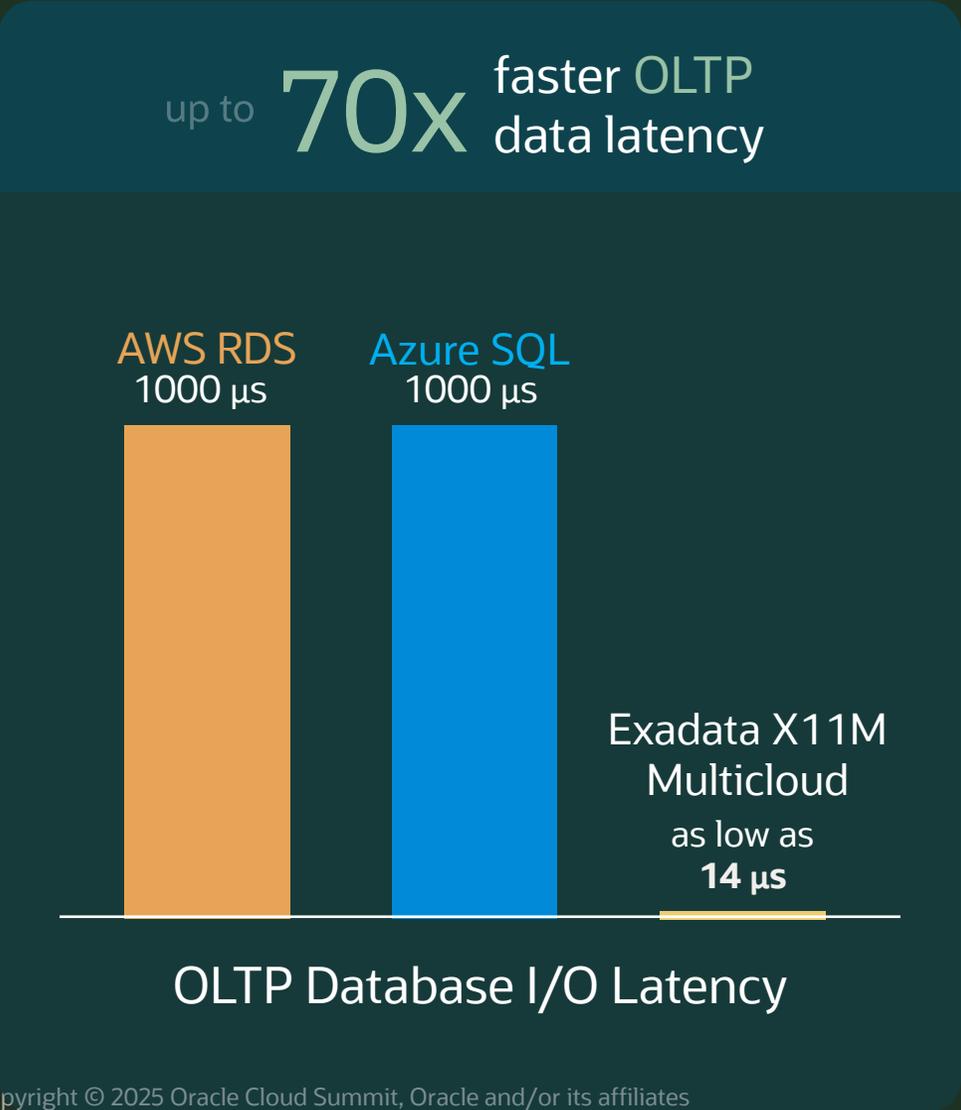
Industry Standard OLTP Benchmark

up to **7.3x** Faster analytic query throughput

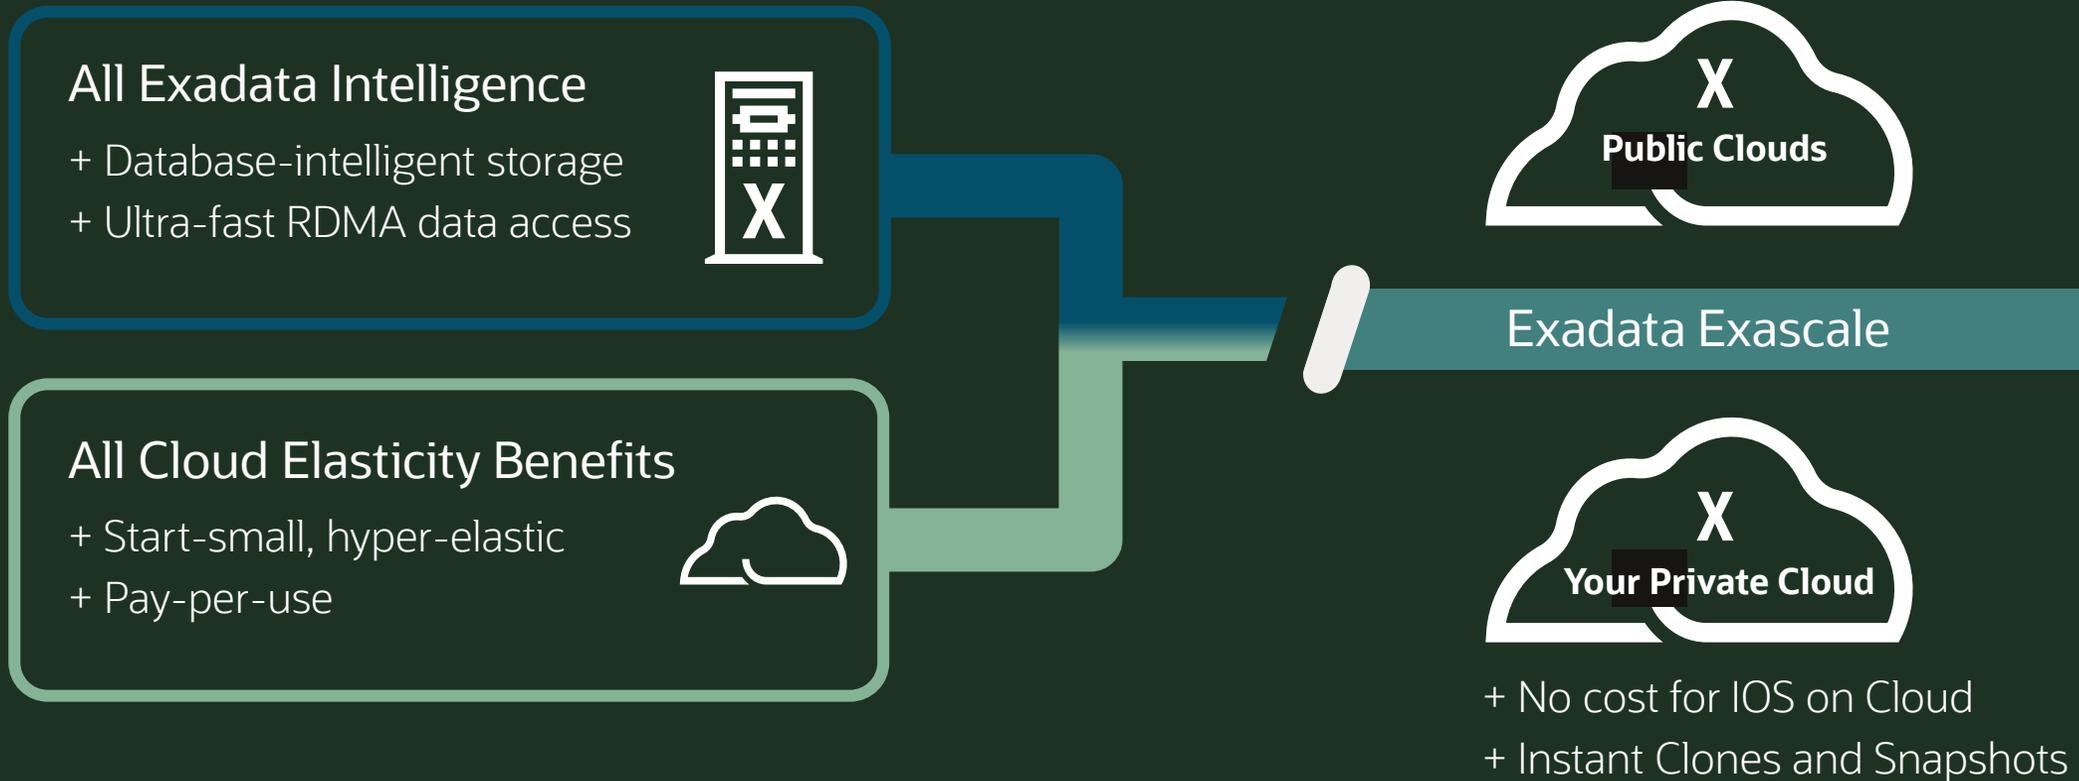


Industry Standard 1TB Benchmark

# Exadata provides exceptional OLTP data latency and analytic throughput



# Exadata Exascale – next generation Exadata software architecture



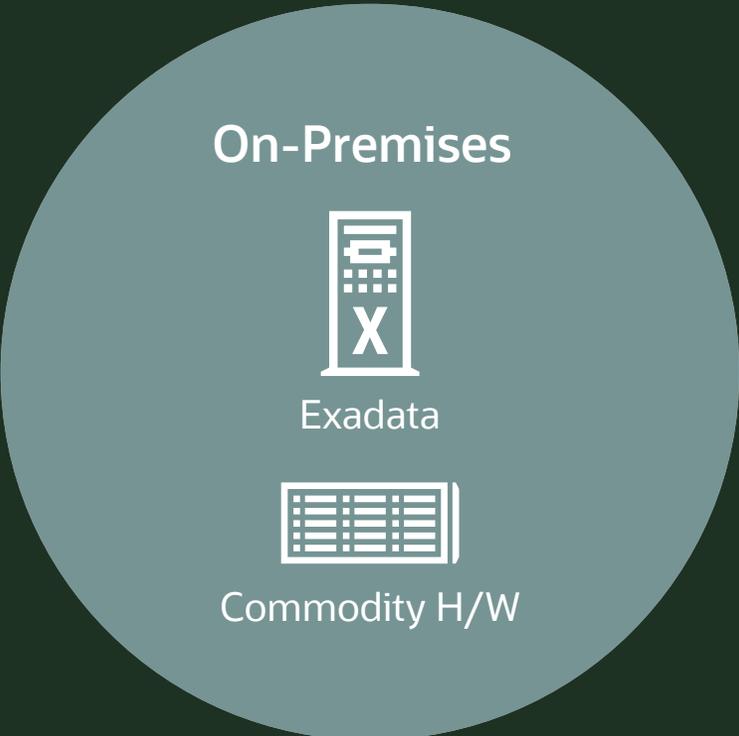
**Best of Exadata Plus Best of Cloud** Engineered to Work Together

# Oracle's Database Strategy

Oracle Database, Exadata, and  
Database Cloud run everywhere



# In your data center



Exadata  
Hardware of your choice

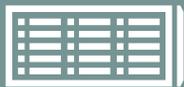


# In all leading public clouds

## On-Premises



Exadata



Commodity H/W

## Cloud

ORACLE  
Cloud

OCI



Google



Azure



Amazon

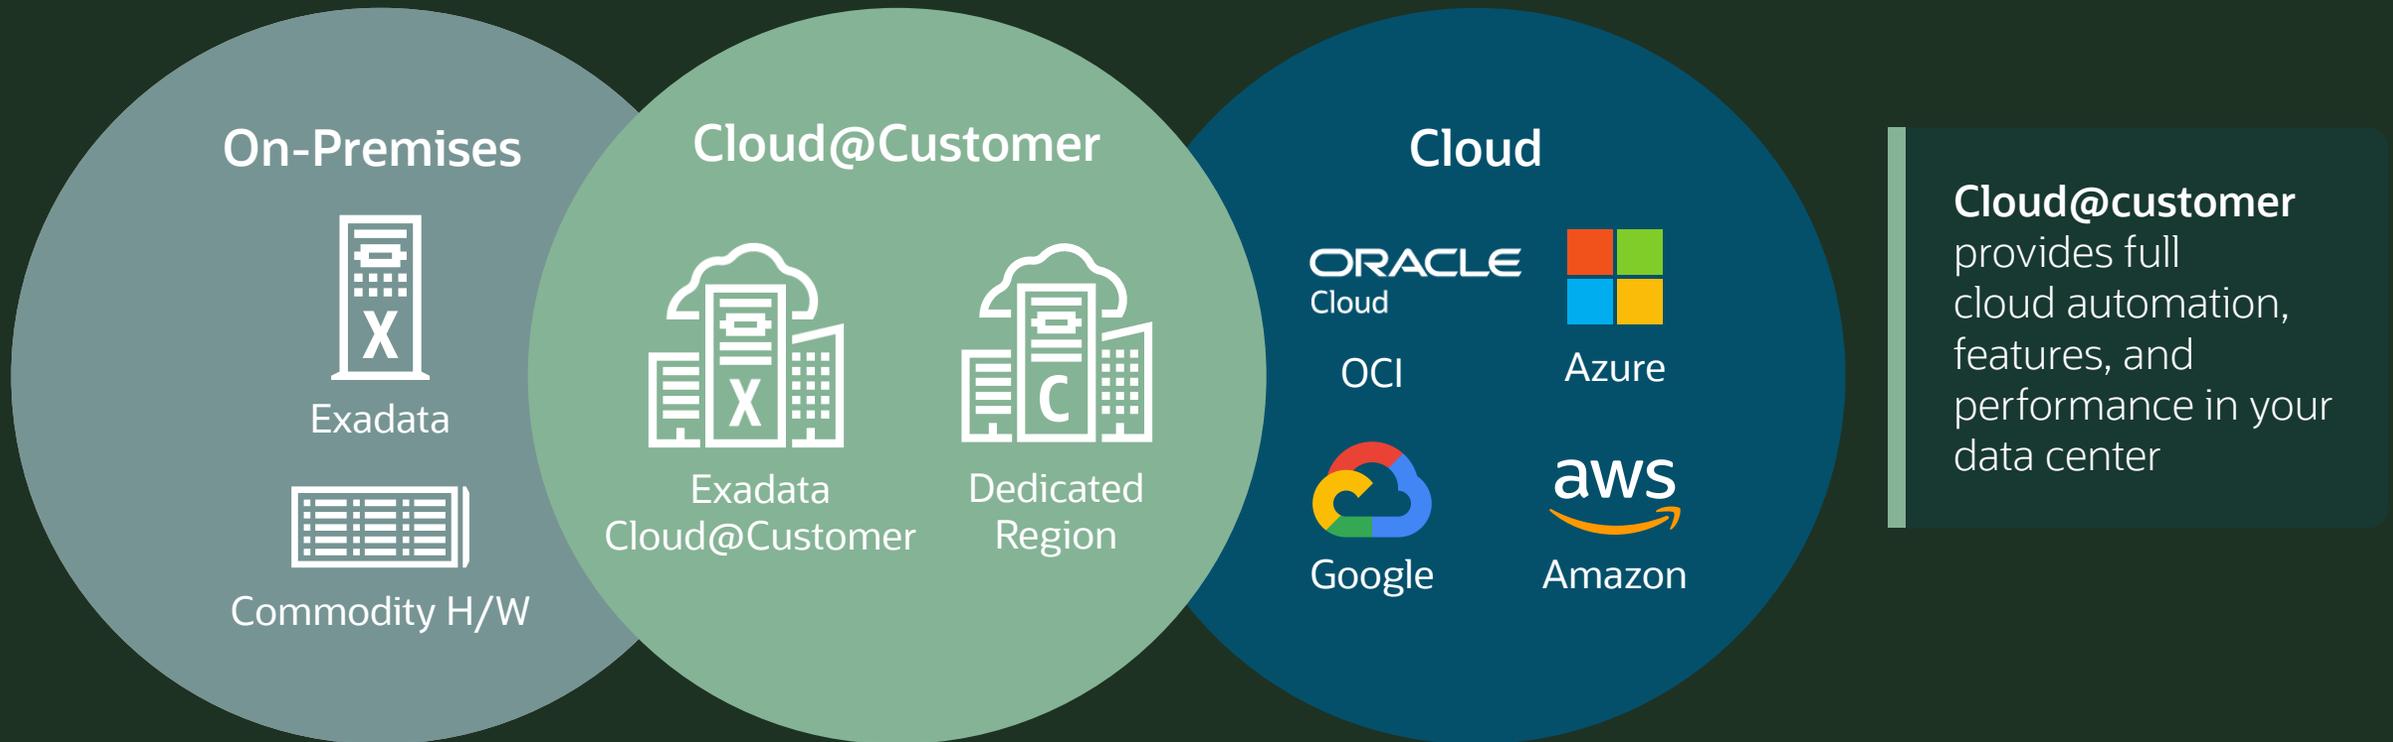
## Oracle Database Multicloud

Full Exadata Cloud automation,  
features, and performance

Use existing Cloud  
Vendor credits

Sub-millisecond latency  
from app to database

# In your data center as a cloud



**Cloud@customer** provides full cloud automation, features, and performance in your data center

Apps and data remain in your data center



# Oracle's Database Strategy

Engineer a completely automated  
Data Platform that delivers the best of  
OLTP, Warehouse, and Data Lake



# Oracle Autonomous Database



Oracle Converged Data Architecture on Exadata provided as a fully automated database service

Any app at any scale transparently gets extreme availability with no effort at low cost

Available in all leading public clouds and on-premises

Autonomous  
Database powers  
1000s of global  
enterprises  
today

accenture

lyft

AON

FedEx®

GENERALI

Premier  
League

experian™

NEC



THOMSON REUTERS



SIEMENS

vodafone

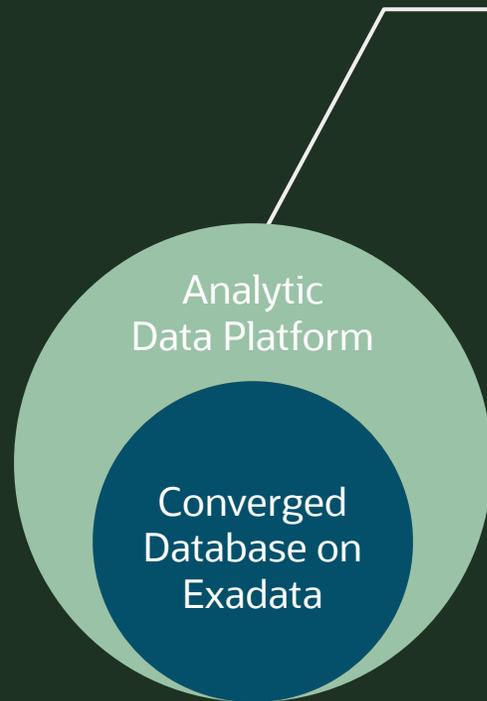
Schneider  
Electric

MARS  
Veterinary Health

SAILGP

xerox

# Oracle Autonomous Database is a complete and simple analytic data platform



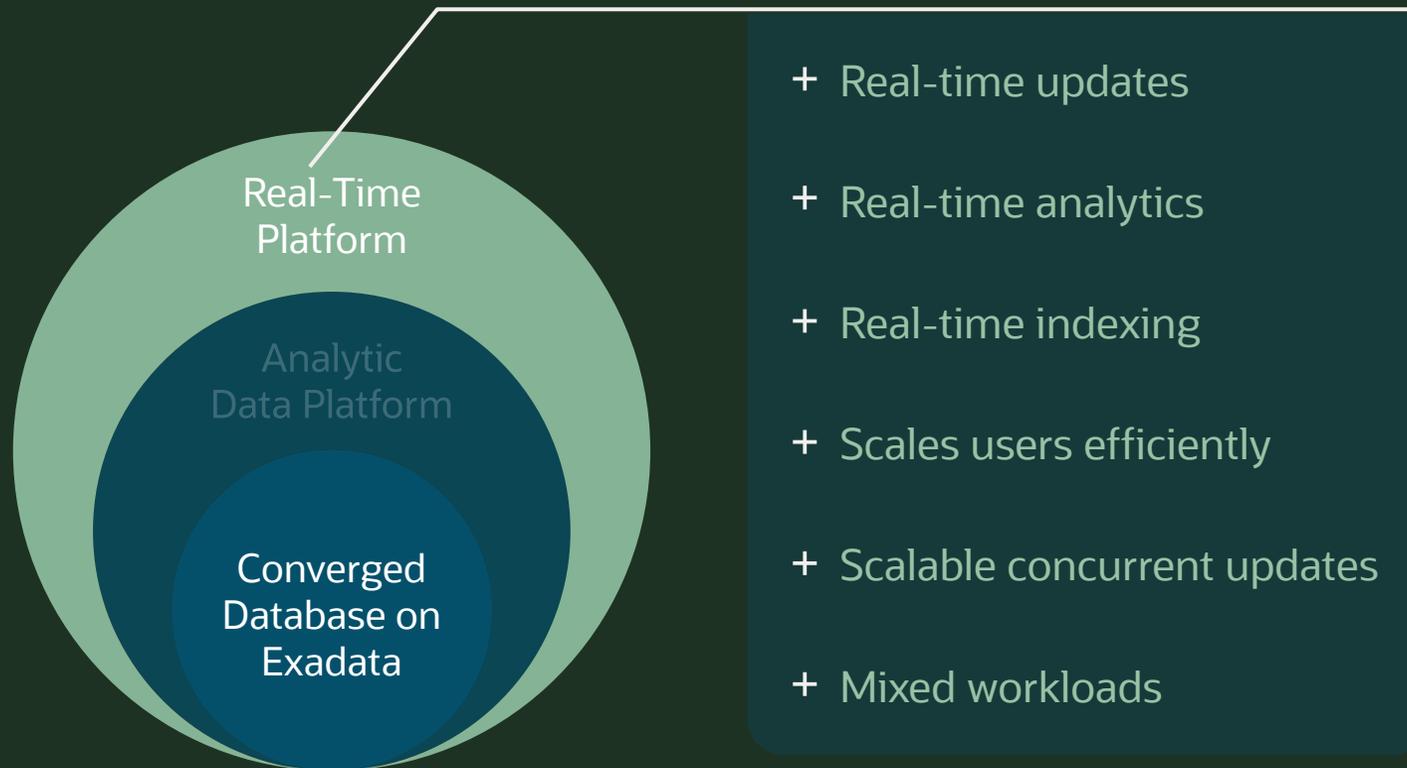
- + Data Warehouse
- + Data Lake
- + Data Sharing
- + Data Engineering
- + Data Catalog
- + Built-in AI/ML

Highly secure

Data can be kept in extreme performance Exadata storage or kept in Iceberg format

All **Data Platform Technologies** Engineered to Work Together

# Autonomous Database is a unique Translytical Data Platform Combines the best of OLTP and Analytics



What competitors describe as their vision of the future is available and mature **now** in Autonomous DB

**Real-Time and Analytics** Engineered to Work Together

# Oracle Translytical Data Platform is rated as having the highest scores in both the current offering and strategy categories

## Forrester Wave: Translytical Data Platforms Q4 2024

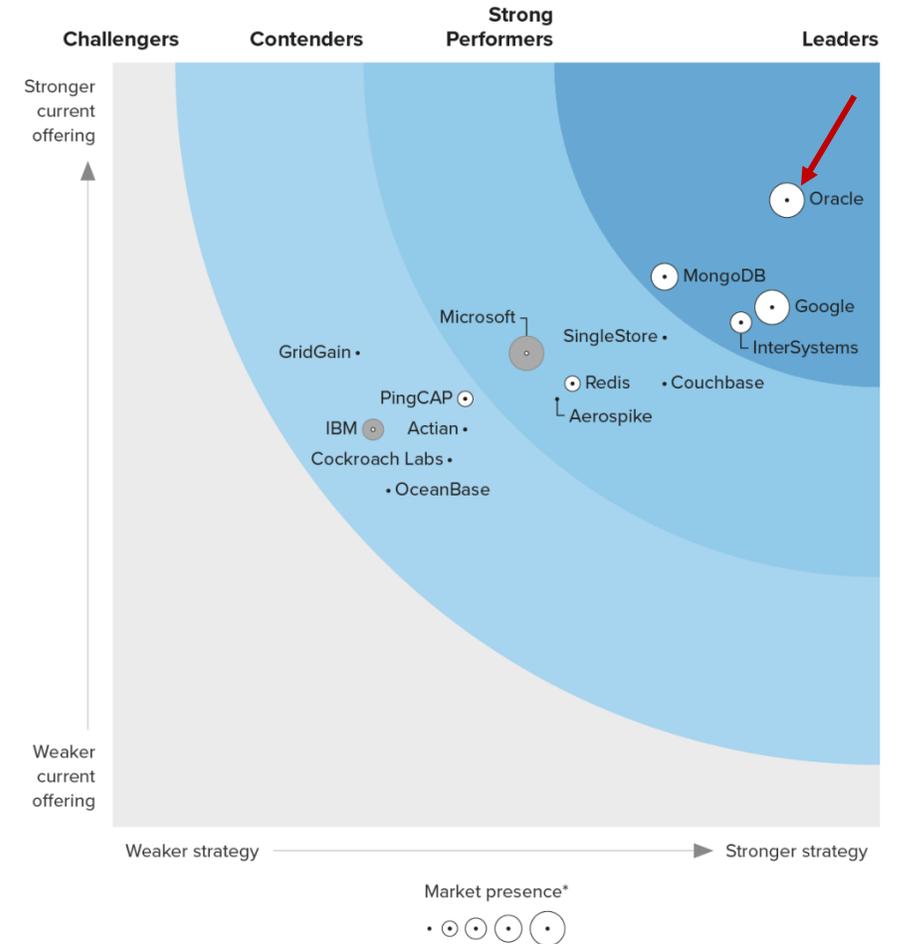
Authored by Noel Yuhanna  
VP, Principal Analyst, Forrester

Source (link forthcoming): <https://reprint.forrester.com/reports/the-forrester-wavetm-translytical-data-platforms-q4-2024-90b043e6/index.html>

### THE FORRESTER WAVE™

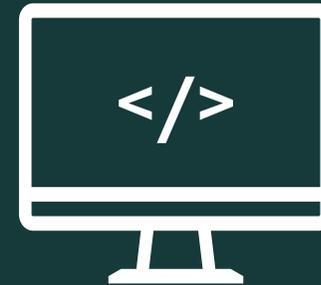
Translytical Data Platforms

Q4 2024



# Oracle's Database Strategy

Engineer apps and data together



# For decades, data professionals have preferred to use relational tables for storing data

**PURCHASE ORDER**

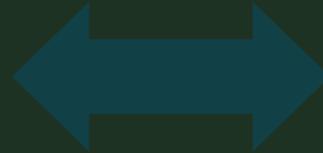


P.O. NUMBER	DATE
23781	01/20/2025

**CUSTOMER**

NAME	PHONE
JOHN SMITH	517-367-7010
COMPANY NAME	EMAIL
SUPREMO CORPORATION	johnsmith@mysmallbiz.org
ADDRESS	DELIVERY INSTRUCTIONS
16050 S. US 22 POLES, MICHIGAN, 43509	

CODE	QUANTITY	PRODUCT DESCRIPTION	UNIT PRICE	AMOUNT
304-98632	4	Brake Discs, Pads & Calipers	111.36	445.44
501-35587	2	Control Arm	60.93	121.86
886-19386	2	Suspension Lift Kit	399.83	799.66



ORDER INFO			
PONUM	DATE	C_ID	INST
...	...	...	...
...	...	...	...
...	...	...	...

CUSTOMER INFO			
C_ID	NAME	PHONE	ADDR
...	...	...	...
...	...	...	...
...	...	...	...

LINE_ITEMS			
PONUM	L_ID	PCODE	QTY
...	...	...	...
...	...	...	...
...	...	...	...

PRODUCT INFO			
PCODE	PNAME	DESC	PRICE
...	...	...	...
...	...	...	...
...	...	...	...



# Relational decomposes application data into its independent reusable components and stores them as rows in separate tables

**PURCHASE ORDER**

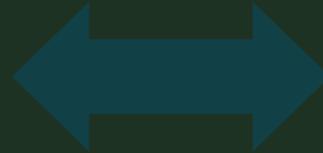
 **SUPREMO CORPORATION**

P.O. NUMBER	DATE
23781	01/20/2025

**CUSTOMER**

NAME	PHONE
JOHN SMITH	517-367-7010
COMPANY NAME	EMAIL
SUPREMO CORPORATION	johnsmith@mysmallbiz.org
ADDRESS	DELIVERY INSTRUCTIONS
16050 S. US 22 POLES, MICHIGAN, 43509	

CODE	QUANTITY	PRODUCT DESCRIPTION	UNIT PRICE	AMOUNT
304-98632	4	Brake Discs, Pads & Calipers	111.36	445.44
501-35587	2	Control Arm	60.93	121.86
886-19386	2	Suspension Lift Kit	399.83	799.66



**ORDER DATA**

PONUM	DATE	C_ID	INST
...	...	...	...
...	...	...	...
...	...	...	...

**CUSTOMER DATA**

C_ID	NAME	PHONE	ADDR
...	...	...	...
...	...	...	...
...	...	...	...

**LINE\_ITEMS DATA**

PONUM	L_ID	PCODE	QTY
...	...	...	...
...	...	...	...
...	...	...	...

**PRODUCT DATA**

PCODE	PNAME	DESC	PRICE
...	...	...	...
...	...	...	...
...	...	...	...



# Sometimes app developers prefer to store data as JSON documents

**PURCHASE ORDER**

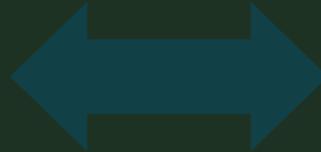
 **SUPREMO CORPORATION**

P.O. NUMBER	DATE
23781	01/20/2025

**CUSTOMER**

NAME	PHONE
JOHN SMITH	517-367-7010
COMPANY NAME	EMAIL
SUPREMO CORPORATION	johnsmith@mysmallbiz.org
ADDRESS	DELIVERY INSTRUCTIONS
16050 S. US 22 POLES, MICHIGAN, 43509	

CODE	QUANTITY	PRODUCT DESCRIPTION	UNIT PRICE	AMOUNT
304-98632	4	Brake Discs, Pads & Calipers	111.36	445.44
501-35587	2	Control Arm	60.93	121.86
886-19386	2	Suspension Lift Kit	399.83	799.66



## ORDER RELATED DATA

```
{
  "po number" : 23781,
  "date"      : "01/20/2025",
  "custid"    : 6543,
  "line items" :
  [ { "prodID" : "304-98632",
      "desc"   : "control arm",
      "unitPrice" : "121.86",
      "quantity" : 2
    },
    ...
  ]
}
```

## CUSTOMER RELATED DATA

```
{
  "csutId" : 6543
  "custName" : "John Smith",
  "address" : "16040 S. US 27 ...",
  "phone" : "517-367-7010",
  "email" : "johnsmith@redline.com"
}
```



# But storing logically independent data in a single document can compromise data consistency, reusability, and queryability

**PURCHASE ORDER**

 **SUPREMO CORPORATION**

P.O. NUMBER	DATE
23781	01/20/2025

**CUSTOMER**

NAME	PHONE
JOHN SMITH	517-367-7010
COMPANY NAME	EMAIL
SUPREMO CORPORATION	johnsmith@mysmallbiz.org
ADDRESS	DELIVERY INSTRUCTIONS
16050 S. US 22 POLES, MICHIGAN, 43509	

CODE	QUANTITY	PRODUCT DESCRIPTION	UNIT PRICE	AMOUNT
304-98632	4	Brake Discs, Pads & Calipers	111.36	445.44
501-35587	2	Control Arm	60.93	121.86
886-19386	2	Suspension Lift Kit	399.83	799.66



```
ORDER RELATED DATA
{
  "po number" : 23781,
  "date"      : "01/20/2025",
  "custid"    : 6543,
  "line items" :
  [ { "prodID" : "304-98632",
      "desc"   : "control arm",
      "unitPrice" : "121.86"
      "quantity" : 2
    },
    ...
  ]
}
```

```
CUSTOMER RELATED DATA
{
  "csutId" : 6543
  "custName" : "John Smith",
  "address" : "16040 S. US 27 ...",
  "phone" : "517-367-7010",
  "email" : "johnsmith@redline.com"
}
```



# Breakthrough insight: the purchase order data being stored is the same, it is just stored in different formats

ORDER DATA			
PONUM	DATE	C_ID	INST
...	...	...	...
...	...	...	...
...	...	...	...

CUSTOMER DATA			
C_ID	NAME	PHONE	ADDR
...	...	...	...
...	...	...	...
...	...	...	...

LINE_ITEMS DATA			
PONUM	L_ID	PCODE	QTY
...	...	...	...
...	...	...	...
...	...	...	...

PRODUCT DATA			
PCODE	PNAME	DESC	PRICE
...	...	...	...
...	...	...	...
...	...	...	...



```
ORDER RELATED DATA
{
  "po number"      : 23781,
  "date"           : "01/20/2025",
  "custid"         : 6543,
  "line items"    :
  [ { "prodID"      : "304-98632",
      "desc"         : "control arm",
      "unitPrice"    : "121.86"
      "quantity"    : 2
    },
    ...
  ]
}
```

```
CUSTOMER RELATED DATA
{
  "csutId"         : 6543
  "custName"       : "John Smith",
  "address"        : "16040 S. US 27 ...",
  "phone"          : "517-367-7010",
  "email"          : "johnsmith@redline.com"
}
```



# New JSON Duality Views declare how to map data between JSON and tables

ORDER DATA			
PONUM	DATE	C_ID	INST
...	...	...	...
...	...	...	...
...	...	...	...

CUSTOMER DATA			
C_ID	NAME	PHONE	ADDR
...	...	...	...
...	...	...	...
...	...	...	...

LINE_ITEMS DATA			
PONUM	L_ID	PCODE	QTY
...	...	...	...
...	...	...	...
...	...	...	...

PRODUCT DATA			
PCODE	PNAME	DESC	PRICE
...	...	...	...
...	...	...	...
...	...	...	...



*JSON operations converted to reads or writes of tables*

```
ORDER RELATED DATA
{
  "po number"      : 23781,
  "date"           : "01/20/2025",
  "custid"         : 6543,
  "line items"    :
  [ { "prodID"      : "304-98632",
      "desc"         : "control arm",
      "unitPrice"    : "121.86"
      "quantity"    : 2
    },
    ...
  ],
  ...
}
```

```
CUSTOMER RELATED DATA
{
  "csutId"         : 6543
  "custName"       : "John Smith",
  "address"        : "16040 S. US 27 ...",
  "phone"          : "517-367-7010",
  "email"          : "johnsmith@redline.com"
}
```

Apps get the best of JSON and Relational, at the same time, on the same data



# Duality views can even provide each app ALL the data it needs in a single document

**PURCHASE ORDER**

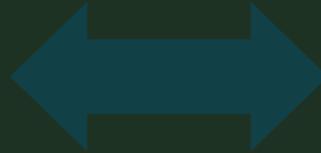
 **SUPREMO CORPORATION**

P.O. NUMBER	DATE
23781	01/20/2025

**CUSTOMER**

NAME	PHONE
JOHN SMITH	517-367-7010
COMPANY NAME	EMAIL
SUPREMO CORPORATION	johnsmith@mysmallbiz.org
ADDRESS	DELIVERY INSTRUCTIONS
16050 S. US 22 POLES, MICHIGAN, 43509	

CODE	QUANTITY	PRODUCT DESCRIPTION	UNIT PRICE	AMOUNT
304-98632	4	Brake Discs, Pads & Calipers	111.36	445.44
501-35587	2	Control Arm	60.93	121.86
886-19386	2	Suspension Lift Kit	399.83	799.66



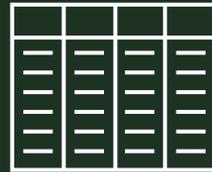
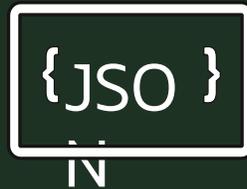
Full Purchase Order  
JSON

```
{
  "PO Number" : "23781",
  "date"      : "01/20/2025",
  "customer"  : "John Smith",
  "Address"   : "16040 S. US 27 ...",
  "line items" :
  [
    {"id" : "304-98632", "desc": ...},
    {"id" : "501-35587", "desc": ...},
    {"id" : "886-19386", "desc": ...}
  ]
}
```

Using JSON with Oracle Database is now simpler than using JSON with a pure JSON databases

# JSON Duality Views together with new Graph Views go beyond converging data types

They **unify** the key data types



ORDERS INFO

CUSTOMER INFO

ORDER ITEM INFO

PRODUCT INFO

**Relational, JSON, and Graph** Engineered to Work Together



*“Oracle’s JSON Relational Duality, a truly revolutionary solution, is perhaps one of the most important innovations in information science in 20 years.”*

Carl Olofson, Research VP,  
Data Management Software, IDC

# Oracle's Database Strategy

Bring AI to your data,  
and engineer them together

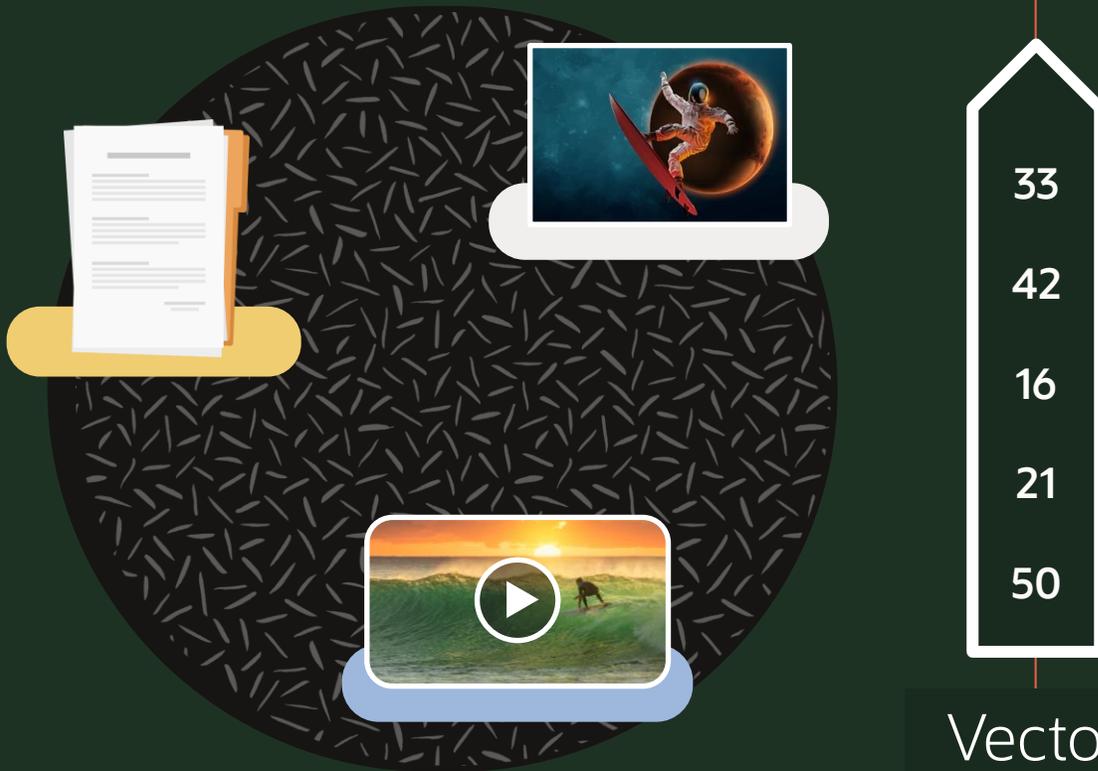




Databases are great at querying business data that is stored as strings, numbers, and dates

But not at querying **human-centric data** such as images, documents, and videos

**AI Vector Search** now makes it easy to query the content of documents, pictures, and videos



An **AI Vector** is a sequence of numbers, called dimensions, that represent the **semantic content** of a document, image, or video

Developers create a vector for an object by just passing the object to a built-in vectorization function

Oracle 23ai natively **stores** vectors and **compares** vectors to find objects with **similar semantic content**

# Engineering AI and data together enables enterprises to easily combine business data search with AI Vector Search

## Example:

Find the top matching support incidents but limit it to incidents for products I own



**Support Rep**  
Jane Doe



- ✓ Spontaneous reboot
- ✓ Resolved
- ✓ Applied OS Update 42

```
SELECT ...  
FROM   Support_Requests S,  
WHERE  S.prod_id IN (SELECT p.id FROM purchases  
                     WHERE cust_id= :me)  
ORDER BY VECTOR_DISTANCE(request_vector, :search_vector);
```

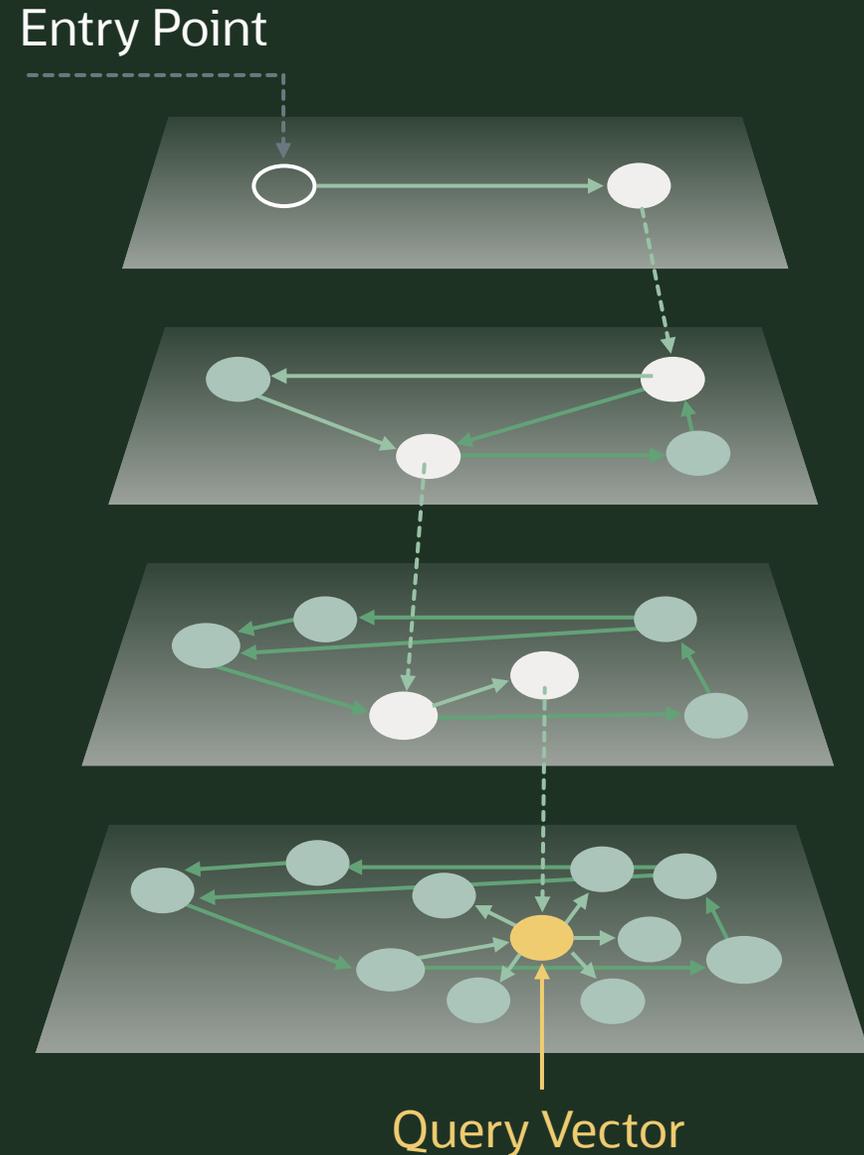
Ultra simple and powerful

Combines customer data, product data, and AI search in 5 lines of SQL!

All data is fully consistent

Single integrated solution

Oracle database accelerates  
AI Vector Search using  
sophisticated new **vector indexes**  
Runs AI Search in milliseconds





Every mission-critical feature of Oracle Database works transparently with AI Vectors

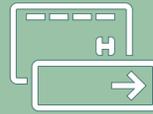
AI Vectors can be used immediately in enterprise apps of any scale or criticality



Real-Application Cluster



Parallel SQL



Transactions



Security



Analytics



Disaster Recovery

# Retrieval Augmented Generation enables users talk to their data



Oracle 23ai improves Generative AI by augmenting LLM prompts with private database content that is found using any combination of data and AI Vector search

Enables LLMs to use business data to produce better and more accurate answers to user questions while keeping business data secure

## Retrieval Augmented Generation (RAG)

# Enterprises are already using Oracle AI Vector Search to create innovative new solutions



## Visual Search for Products

Find products that are similar to a user provided image



## Real-time offer management

Enable merchants to present the right offers to consumers at checkout



## Fraud Detection

Multi-lingual search for similar insurance claims that were found to be fraudulent



## Real-time intelligent assistant

Use RAG to answers customer questions about products

# Database 23ai includes many other AI Capabilities



Translating natural language into SQL using **Select AI**



Implementing over **30 ML algorithms** inside the DB



Running Distributed AI with **GoldenGate 23ai**



Scaling Mission-critical AI using **Exadata**



Creating vectors inside DB using the **ONNX** framework



Interfacing with popular AI Dev Tools incl. **LangChain**

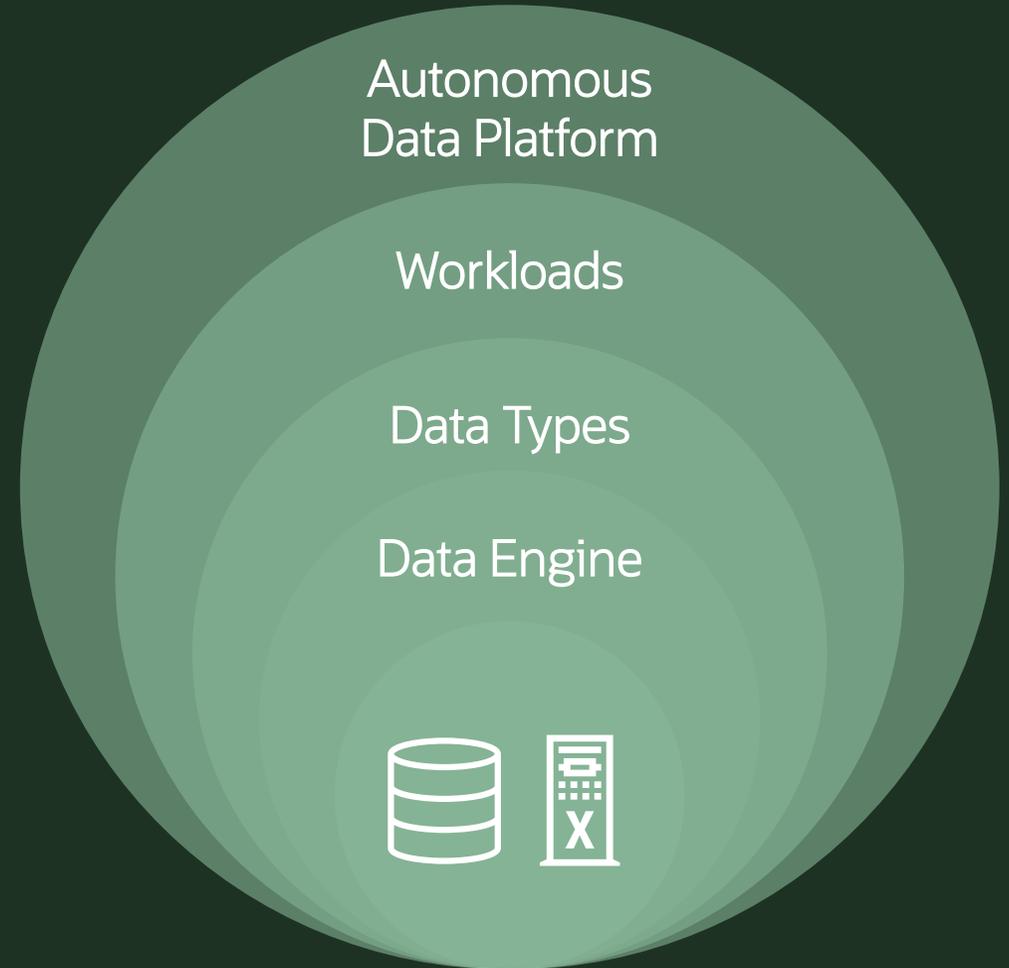
All Oracle Database AI capabilities are provided for no additional charge

# Key Takeaways

Complexity is **unavoidable** when apps are built using technologies designed in isolation

**Oracle's strategy** is to engineer all modern data technologies to work together

Improving synergy, efficiency, and data integrity while helping reduce costs



# Engineered to Work Together



All data types  
and workloads



Hardware and  
software



OLTP, DW,  
and Data Lake



Full mission  
critical stack



Apps and  
data together



AI and  
app dev

ORACLE